right control information from the tag is described as follows: instead of being located at different positions from stage to stage, the two-bit in-band control signal should be always at the fixed position, say, the first two bits of the tag, such that the control circuitry at each stage can always read the leading two bits of the routing tag to make the routing decision. To achieve this, when a packet reached the output port of a stage and before entering the next stage, the second bit of the routing tag is shifted to the end of the tag, or just removed from the tag, by a simple dedicated 1×1 switching circuitry which is appended to every output port. In other words, each stage here actually performs the routing of the packet and the re-generation of the routing tag for the next stage. In this way, the first two bits are $1d_{\gamma(1)}$ when entering stage 1, and $1d_{\gamma(2)}$ when entering stage 2, and so on, that is, the leading two bits of the routing tag of the packet entering each stage j are always $1d_{\gamma(j)}$, the right control signal required by the control circuitry of that stage. As a consequence, the control circuitries can be identical at all stages.--;

Page 177, line 7 through page 178, line 17, replace the Paragraph as follows:

--Example 1. To demonstrate this generalized self-routing mechanism, consider network **2900** of FIG. 29. The destination address binary($d_1d_2d_3d_4$) for a packet is 1110. The guide has been computed earlier as the sequence 2, 4, 1, 3. Thus, $d_{\gamma(1)} = d_2 = 1$, $d_{\gamma(2)} = d_4 = 0$, $d_{\gamma(3)} = d_1 = 1$, and $d_{\gamma(4)} = d_3 = 1$, so the data packet is prepended with the binary stream $1d_{\gamma(1)}d_{\gamma(2)}d_{\gamma(3)}d_{\gamma(4)} = 11011$ as the routing tag. Each cell in the network is a sorting cell with respect to the linear order of

$$10\ (\text{'0-bound'}) \prec 00\ (\text{'idle'}) \prec 11\ (\text{'1-bound'}).$$

Recall that such a routing cell always routes 0-bound signal (with control bits 10) to

output 0 and 1-bound signal (with control bits 11) to coutput-1 when there is no output contention. Therefore, assuming no output contention occurs at each of the nodes along the path, upon entering the first stage at routing cell **2910**, the two leading control bits, namely, 11, are used to set the connection state of the cell **2910** to "cross" in this case since the signal enters the routing cell from its upper input, resulting in routing the packet to the lower output of the cell, that is, to the output address 1101 at that stage. Meanwhile the second bit of the in-band control signal, namely 1, is consumed by the appended 1×1 device (omitted in the drawing) and thus the new in-band control signal to the next stage becomes 10. Next, exchange $X_{(3\ 4)}$ leads the packet from the output address 1101 of stage 1 to the input address 1110 of stage 2. Then the new in-band control signal, namely 10, is used to set the stage-2 cell **2920** to the "bar" state, resulting in routing to output address 1110. Meanwhile the second bit of the in-band control signal, namely 0, is again consumed and thus the new in-band control signal to the next stage (stage 3) becomes 11. Next, exchange $X_{(1\ 4)}$ leads the packet from the output address 1110 of stage 2 to the input address 0111 of stage 3. Then the new 2-bit control sequence, namely 11, are used to set cell **2930** to the bar state, resulting in routing the packet to the output address 0111. Then the second bit of the in-band control signal, namely 1, is again consumed before entering stage 4. Finally, exchange $X_{(2\ 4)}$ leads the packet from the output address 0111 of stage 3 to the input address 0111 of stage 4. The remaining two control bits, namely 11, is used to set the cell **2940** to the bar state, then the packet is routed to the output address 0111, and finally led to its desired destination address 1110 through the output exchange $X_{(4\ 3\ 2\ 1)}$.---.

Page 179, line 14 through page 180, line 6, replace the Paragraph as follows:

-- As already mentioned in the Background Section, and now well understood because of the foregoing description, the main reason behind the trial-and-error procedure of prior art was that such techniques had not had the benefit of a fundamental theoretical approach of determining the routing tag $d_{\gamma(1)}d_{\gamma(2)}\ldots d_{\gamma(n)}$ or $1d_{\gamma(1)}d_{\gamma(2)}\ldots d_{\gamma(n)}$ from the guide of a bit-permuting network. The routing tag for the particular $2^n\times 2^n$ networks studied in the prior art is the destination address $d_1d_2\ldots d_n$ of a packet plus possibly an activity bit up front. By happenstance, the general routing tag $d_{\gamma(1)}d_{\gamma(2)}\ldots d_{\gamma(n)}$ coincides with the destination address $d_1d_2\ldots d_n$ in the special case when the guide of a banyan-type network is the monotonically increasing sequence (i.e., the sequence 1, 2, ... , n). As is now readily deduced, the destination address can be used as the routing tag only for those $2^n\times 2^n$ banyan-type networks with monotonically increasing guide. --.

Page 195, line 9 through page 199, line 6, replace the Paragraphs as follows:

-- One of the criteria mentioned in the above in choosing the proper switch to fill the dilated node in a b-line version of a banyan-type network is a "partial property" of being nonblocking. Explicitly this partial property means the guarantee to route the maximum possible number of 0-bound signals to the 0-output group and the maximum possible number of 1-bound signals to the 1-output group. For a 2b-to-b concentrator composed of interconnected routing cells (plus possibly 1×1 elements), the nature of a concentrator in routing the smallest m–n signals to the 0-output group and the largest n signals to the 1-output group is precisely equivalent to this guarantee. Therefore, a 2b-to-b concentrator composed of interconnected routing cells meets this criterion perfectly for filling the

dilated node in a b-line version of a banyan-type network.

The other criterion in choosing the proper switch to fill the dilated node in a b-line version of a banyan-type network is the compatibility with self-routing over the banyan-type network. The 2b-to-b concentrator composed of interconnected routing cells again meets the criterion perfectly. As a switch constructed by a partial sorting network, a concentrator possess a natural self-routing mechanism. When the 2b-to-b concentrator fills every dilated node of the b-line version of the banyan-type network, the whole network becomes a large multi-stage interconnection network of routing cells. The marriage between the self-routing mechanism over the partial sorting networks with the self-routing mechanism over the banyan-type network, as to be detailed in the next sub-section, creates a self-routing mechanism over the said large multi-stage interconnection network of sorting cells.

C4 Cont

Remark. As before, if idle expressions are disallowed in the system, the 2b-to-b concentrator composed of interconnected routing cells can be substituted by a 2b-to-b concentrator composed of interconnected 0-1 sorting cells. The same applies throughout the next sub-section.

6. Self-routing over a multi-stage interconnection network of concentrators

Hereafter unless otherwise specified, all concentrators refer to those constructed by partial sorting networks.

Recall the classification of multi-stage networks of sorting cells into

routing networks and partial sorting networks. The in-band control signal of a packet is preserved through a partial sorting network. On the other hand, it changes from stage to stage when the packet traverses a routing network, e.g., a banyan-type network under basic self-routing control. The b-line version of a $2^n \times 2^n$ banyan-type network is a hybrid between a routing network and a partial sorting network when every dilated node in it is filled with a 2b-to-b concentrator composed of interconnected routing cells. The hybrid network may be viewed as composed of n "super stage" of concentrators. At each super stage, a packet traverses through a partial sorting network, which is by itself a multi-stage network of routing cells, and the in-band control signal of a packet changes only between super-stages.

C4 cont

The $b2^n$ outputs of the hybrid network are in 2^n groups of the size b. The destination of a packet is an output group rather than an individual output in an output group. In accordance with the present invention, upon entering a generic $2^n \times 2^n$ banyan-type network with the guide $\gamma(1), \gamma(2), \ldots, \gamma(n)$, a packet destined for the output at the address $d_1 d_2 \ldots d_n$ is preceded by the routing tag $1 d_{\gamma(1)} d_{\gamma(2)} \ldots d_{\gamma(n)}$ and the in-band control signal to stage-j switching cell is $1 d_{\gamma(j)}$. The same routing tag still applies in the b-line version of the banyan-type network in which every dilated node is filled by a 2b-to-b concentrator when the packet is destined for the output group at the address $d_1 d_2 \ldots d_n$, and, for $1 \leq j \leq n$, the in-band control signal to a concentrator in the j^{th} super-stage is $1 d_{\gamma(j)}$. More explicitly, the in-band control signal to every routing cell in a concentrator at the j^{th} super-stage is $1 d_{\gamma(j)}$. As the packet progressed through the hybrid network composed of many stages of routing cells, the in-band control signal to a routing cell changes only

upon the exit from a concentrator. That is, the bit $d_{\gamma(j)}$ is consumed not by any generic routing cell inside a concentrator at the j^{th} super-stage but rather by certain extra circuitry installed at the output end of the concentrator. This extra circuitry handles each packet separately and hence consists of 2b parallel 1×1 switching elements. There may exist other 1×1 elements in the 2b-to-b concentrator, e.g., delay elements in maintaining the synchronization across the stage and annihilators of misrouted packets.

Example 7. The guide of the 16×16 divide-and-conquer network is the sequence 1, 2, 3, 4. The network **6900** shown in FIG. 69 is the 8-line version of the 16×16 divide-and-conquer network. This is a 128×128 network, and each of the dilated nodes is 16×16. Thus fill every dilated nodes (e.g. **6901**) with a 16-to-8 concentrator consists of multi-stage interconnected routing cells plus 1×1 elements. The 128 outputs of this network are partitioned into 16 output groups of the size 8. Each output group is associated with a 4-bit address. A packet is destined for an output group rather than a specific output in the group. That is, the routing of a signal to any port within a group is just as good as routing to any other port in the group. When the destined output group is at the address $d_1d_2d_3d_4$, the initial routing tag of the packet is $1d_{\gamma(1)}d_{\gamma(2)}d_{\gamma(3)}d_{\gamma(4)} = 1d_1d_2d_3d_4$. The in-band control of the packet to every routing cell in the concentrator at the 1^{st} super-stage is $1d_1$. Upon exiting that concentrator, the bit d_1 in the routing tag is consumed by a 1×1 element in the concentrator. Thus the routing tag upon entering the 2^{nd} super-stage is $1d_2d_3d_4$. And so on.--.

C4 Cont

Page 200, line 8 through page 200, line 16, replace the Paragraph as follows:

--A concentrator composed of interconnected routing cells is a point-to-point switch that

C5

routes 0-bound, 1-bound, and idle packets to 0- and 1-output groups; it satisfies the desirable characteristic of always routing the maximum possible number of 0-bound (resp. 1-bound) signals to its 0-output group (resp. 1-output group). For a multicast switch that routes 0-bound, 1-bound, idle, and bicast packets to 0- and 1-output groups, a corresponding desirable characteristic is to route the maximum total number of 0-bound and bicast signals to the 0-output group and the maximum total number of 1-bound and bicast signals to the 1-output group. This concept is formulated in the next definition.--.

Page 206, line 8 through page 206, line 15, replace the Paragraph as follows:

-- Example 10. A generic binary address of a $2^6 \times 2^6$ banyan-type network is $b_1b_2b_3b_4b_5b_6$. The entirety of 2^6 output addresses is a 6-dimensional binary cube $S_1 \times S_2 \times \ldots \times S_6$, where each $S_j = \{0, 1\}$ corresponds to the two possible values of b_j. One of the rectangles of this 6-dimentional binary cube can be the subset in the form of $\{0, 1\} \times \{0\} \times \{0, 1\} \times \{1\} \times \{0, 1\} \times \{1\}$, which contains 2^3 output addresses, namely, 000101, 000111, 001101, 001111, 100101, 100111, 101101, and 101111, so this is a 3-dimentional rectangle. The number of 3-dimensional rectangles in the 6-dimensional binary cube is $2^{6-3} *_6C_3 = 8*(6*5*4)/(3*2) = 160$.--.

Page 209, line 18 through page 210, line 8, replace the Paragraph as follows:

-- The bicast cell can be modified for the priority treatment similarly as before. The primary in-band control signal used at each stage j is still $Q_{Y(j)}$, while the priority code $p_1 \ldots p_r$ serves as the tiebreaker when the two packets arrived at the same cell are both 0-

bound or both 1-bound. The switching control at each stage consumes the leading quaternary symbol (or rotated it to the end of the routing tag) and rotates the priority code to the position behind the next quaternary symbol. Therefore, the underlying methodology for the realization of this (multicast) self-routing mechanism over a banyan-type network and the implementation of the related circuitry is very similar to the case of basic (point-to-point) self-routing mechanism employed in banyan-type network.--.

Page 212, line 8 through page 212, line 18, replace the Paragraph as follows:

--Similar to the case of self-routing over a multi-stage interconnection network of concentrators, when the underlying banyan-type network of a multi-stage interconnection network of multicast concentrators is replaced by a more general bit-permuting network, the self-routing control mechanism still applies. More precisely, when the replacing bit-permuting network is a $2^n \times 2^n$ k-stage bit-permuting network with the guide $\gamma(1), \gamma(2), \ldots, \gamma(k)$, where γ is a mapping from the set $\{1, 2, \ldots, k\}$ to the set $\{1, 2, \ldots, n\}$, a packet destined for output groups with the rectangular set of addresses encoded by $Q_1, Q_2, \ldots, Q_n$ is prefixed with the routing tag $Q_{\gamma(1)}Q_{\gamma(2)}\ldots Q_{\gamma(k)}$. For $1 \leq j \leq k$, the in-band control signal to a multicast concentrator in the j^{th} super-stage is $Q_{\gamma(j)}$, and this quaternary symbol in the routing tag is consumed or rotated to the end of the routing tag by the j^{th} super-stage. Note that if $\gamma(p) = \gamma(q)$ in the guide of the network, where $p<q$, the q-th symbol of the routing tag $Q_{\gamma(q)}$ will repeat the p-th symbol $Q_{\gamma(p)}$, when $Q_{\gamma(p)} = Q_{\gamma(q)}$ = 'bicast', the packet may be bicasted at stage-p and then be bicasted again at stage-q such that undesired extra copies of the packet will be produced. Therefore, whenever $\gamma(p) = \gamma(q)$ in the guide of the network, the bicasting function of the whole stage of switching nodes at either stage-p or

stage-q should be disabled to prevent such situation. The remaining parts of the control coincide with the above.--.

Page 225, line 1 through page 235, line 10, replace the Paragraphs as follows:

--J. SELF-ROUTING MULTICASTING TO AN ARBIRTRARY SET OF OUTPUT ADDRESSES OVER A BANYAN-TYPE NETWORK

1. Multicasting to an arbitrary set of output addresses over a banyan-type network

Consider a $2^n \times 2^n$ multicast switch allows a packet to be destined for an arbitrary subset of the 2^n output addresses. A straightforward encoding scheme by 2^n bits is to use one bit for each output to indicate whether this output is included in the set of destination addresses. In general, the number of bits required in encoding an arbitrary set of destination addresses is at least 2^n. The switching control mechanism has to read all the 2^n bits before making the decision and thus incurs $\Omega(2^n)$ latency. If a self-routing switch were to implement this multicast function, then every primitive switching element in it would have to buffer 2^n bits and thereby incurs $\Omega(2^n)$ complexity in both the hardware and the latency.

The present disclosure presents a self-routing mechanism for multicasting packets to an arbitrary set of output addresses over a generic $2^n \times 2^n$ banyan-type network, where the destination set is encoded in 2^n-1 quaternary symbols, the hardware complexity of a stage-j 2×2 switching element is O(j), and the switching delay through the element is O(1).

Definition J1: "guiding sequence of a packet". For a $2^n \times 2^n$ banyan-type

network with the guide $\gamma(1), \gamma(2), \ldots, \gamma(n)$ where γ is a permutation of the integers from 1 to n, if the destination addresses of a packet include the address $d_1d_2\ldots d_n$, then $d_{\gamma(1)}d_{\gamma(2)}\ldots d_{\gamma(n)}$ is said to be a "guiding sequence" of that packet.

For example, for a $2^6\times2^6$ banyan-type network with the guide being 5, 4, 6, 1, 3, 2, if the destination addresses of a multicast packet in this network comprise 001010 (address 1), 011001 (address 2) and 110101 (address 3), for address 1 where $d_1d_2d_3d_4d_5d_6 = 001010$, that is, $d_1=0$, $d_2=0$, $d_3=1$, $d_4=0$, $d_5=1$, and $d_6=0$, then $d_{\gamma(1)}d_{\gamma(2)}d_{\gamma(3)}d_{\gamma(4)}d_{\gamma(5)}d_{\gamma(6)} = d_5d_4d_6d_1d_3d_2 = 100010$ is a guiding sequence of this packet; for address 2 where $d_1d_2d_3d_4d_5d_6 = 011001$, $d_{\gamma(1)}d_{\gamma(2)}d_{\gamma(3)}d_{\gamma(4)}d_{\gamma(5)}d_{\gamma(6)} = d_5d_4d_6d_1d_3d_2 = 001011$ is also a guiding sequence of this packet; for address 3 where $d_1d_2d_3d_4d_5d_6 = 110101$, $d_5d_4d_6d_1d_3d_2 = 011101$ is another guiding sequence of this packet.

In accordance with the present invention, the $2^n\times2^n$ multicast switch is constructed from a $2^n\times2^n$ banyan-type network of bicast cells. A coding scheme that uses 2^n-1 quaternary symbols (equivalent to $2^{n+1}-2$ bits) is adopted to encode the arbitrary set of the destination addresses of a multicast packet. As before, the four values of a generic quaternary symbol are '0-bound', '1-bound', 'idle', and 'bicast'. Associated with every binary string S of a length from 0 to n−1 is a quaternary symbol Q_S. The information of the destination addresses of a packet is encoded by a sequence of quaternary symbols. Upon entering the first stage, the packet is prefixed with the following sequence of quaternary symbols as the routing tag

$$Q_{\sim}Q_0Q_1Q_{00}Q_{10}Q_{01}Q_{11}Q_{000}Q_{100}Q_{010}Q_{110}Q_{001}Q_{101}Q_{011}Q_{111}Q_{0000}Q_{1000}\cdots$$

That is, quaternary symbols associated with shorter strings precede those associated with longer strings. Among symbols associated with equally long strings, the

order is right-to-left lexicographical.

Consider a generic $2^n \times 2^n$ banyan-type network with the guide being the sequence $\gamma(1), \gamma(2), \ldots, \gamma(n)$. By definition, $d_{\gamma(1)}d_{\gamma(2)}\ldots d_{\gamma(n)}$ is a guiding sequence of a packet when the destination addresses of that packet include the address $d_1d_2\ldots d_n$. The assignment of value to Q_S pertains to the guiding sequence as follows:

Q_S = 'idle' if S is not a prefix of any guiding sequence, or

Q_S = '0-bound' if S0 is a prefix of a guiding sequence but S1 is not, or

Q_S = '1-bound' if S1 is a prefix of a guiding sequence but S0 is not, or

Q_S = 'bicast' if both S0 and S1 are prefixes of guiding sequences.

The in-band control signal used by the bicast cell at each stage is the leading quaternary symbol in the routing tag of the packet. Recall the control of a bicast cell, if the leading quaternary symbol of one of the two packets arrived at the bicast cell is 'bicast' and that of the other packet is 'idle', then the former packet is bicasted to both outputs of the bicast cell; otherwise, the bicast cell sorts the two packets according to their leading quaternary symbols with respect to the partial order of '0-bound' $\prec$ 'idle' $\prec$ '1-bound' and '0-bound' $\prec$ 'bicast' $\prec$ '1-bound'. This describes the switching control over a single bicast cell. Meanwhile, in accordance with the present invention, there is also the switching control at the network level, which modifies the routing tag of a packet when the packet exits from one stage for the next stage as described below.

A packet routed to output-0 of a stage-1 cell retains only every other quaternary symbol starting with the second symbol in the routing tag. Those non-retained symbols in between are replaced by space fillers in order to maintain the bit pipelining, which assume arbitrary quaternary values. The routing tag thus becomes

$$Q_0\square Q_{00}\square Q_{01}\square Q_{000}\square Q_{010}\square Q_{001}\square Q_{011}\square Q_{0000}\square\ldots$$

where the symbol "□" indicates a space filler. Similarly, a packet routed to output-1 of a stage-1 cell retains only every other quaternary symbol starting with the third symbol in the routing tag. The routing tag thus becomes

$$Q_1\square Q_{10}\square Q_{11}\square Q_{100}\square Q_{110}\square Q_{101}\square Q_{111}\square Q_{1000}\square\ldots$$

A packet routed to output-0 of a stage-j cell retains only every other real quaternary symbol starting with the second real symbol in the routing tag, while a packet routed to output-1 of a stage-j cell retains only every other real quaternary symbol starting with the third real symbol in the routing tag. Note that space fillers are not regarded as real quaternary symbols. Again, space fillers replace those non-retained symbols in order to maintain the bit pipelining. After stage-2, the routing tag becomes one of the following:

$$Q_{00}\square\square\square Q_{000}\square\square\square Q_{001}\square\square\square Q_{0000}\square\square\square\ldots$$

$$Q_{01}\square\square\square Q_{010}\square\square\square Q_{011}\square\square\square Q_{0100}\square\square\square\ldots$$

$$Q_{10}\square\square\square Q_{100}\square\square\square Q_{101}\square\square\square Q_{1000}\square\square\square\ldots$$

$$Q_{11}\square\square\square Q_{110}\square\square\square Q_{111}\square\square\square Q_{1100}\square\square\square\ldots$$

By induction, upon entering stage-j, there are $2^{j-1}-1$ space fillers between

every two consecutive real quaternary symbols. Each stage-j output preserves every 2^j-th symbol starting at a certain real quaternary symbol. A bit-clock counter is required to count from 1 to 2^j in order to implement this operation. The hardware complexity of such a counter is O(j).

Example 1. FIG. 76 illustrates the multicasting of a first packet (**7601**) toward two outputs, namely, 000 and 011, and a second packet (**7602**) toward four outputs, namely, 010, 100, 101, and 111, of an 8×8 banyan network (**7600**). The coding of the destination addresses information as a sequence of quaternary symbols into the routing tag of each packet is as follows. The quaternary symbols 0-bound, 1-bound, idle, and bicast are abbreviated as 0, 1, I, and B, respectively. Since the guide of the 8×8 banyan network, $\gamma(1), \gamma(2), \gamma(3)$, is simply the monotonic sequence 1, 2, 3, the guiding sequence of a packet corresponding to its destination address $d_1d_2d_3$ is $d_{\gamma(1)}d_{\gamma(2)}d_{\gamma(3)} = d_1d_2d_3$. Therefore, the guiding sequences for the first packet are 000 and 011, and those for the second packet are 010, 100, 101, and 111. For the first packet, the first symbol $Q_{\sim}$ in the routing tag is 0 because, according to the rules of the assignment of value to each Q_S as described above, the string S now is a null string, thus since the string S0 = "0" is a prefix of a guiding sequence, that is, 000, of the first packet, but S1 = "1" is not a prefix of any guiding sequences of the first packet, the condition for the case Q_S = '0-bound' is matched; for the second symbol Q_0, S = "0" now and hence both S0 ("00") and S1 ("01") are prefixes of the guiding sequences of the first packet, so Q_0 = B; the third symbol Q_1 = I because S = "1" now which is not a prefix of any guiding sequence of the first packet; and so on. The routing tag for the first packet is thus generated as "0BI0I1I" (**7611**).

Similarly, the routing tag for the second packet is thus generated as "B1BIB01" (**7621**). The first packet enters the bicast cell (**7631**) in the first stage from its input-1, since its leading symbol is "0" and the other input of the cell is idle, the cell sets its connection state to be cross and routes the first packet to its output-0. Meanwhile, the second, fourth, and sixth symbols, that is, "B", "0", and "1", are retained in the routing tag (**7612**) as the first packet exits from the cell **7631**. From now on only the routing tag of the packet is shown in the drawing for simplicity. Then the first packet enters the bicast cell (**7632**) in the second stage with the leading symbol "B". Since the other input of the cell **7632** is idle, the packet is bicasted to both outputs. On exiting from the cell **7632**, the copy of the packet at the output-0 retains every other real quaternary symbol starting with the second real symbol in the routing tag "B□0□1□" (**7612**), "0", and thus gives the new routing tag "0□□□" (**7613**), while the copy of the packet at the output-1 retains every other real quaternary symbol starting with the third real symbol in the routing tag "B□0□1□" (**7612**), "1", and thus gives the new routing tag "1□□□".

Note that since banyan-type networks are blocking, because of the possible packet collision inside the network, there is no guarantee that a packet will reach all or any of the intended destinations. By applying the technique of statistical line grouping as described in Section H, the blocking problem can be alleviated while the self-routing mechanism can still be applied.

2. Construction of self-routing nonblocking switches in both point-to-point and multicast senses

Recall that the recursive plain 2-stage interconnection network of cells associated with the n-leaf rightist tree is the $2^n \times 2^n$ baseline network. The first recursive step in this construction is depicted in FIG. 77A for n=4. Similarly, the recursive 2-swap construction associated with the n-leaf rightist tree is the $2^n \times 2^n$ banyan network. Both networks are examples of recursive bit-permuting 2-stage construction (as defined in Section F) associated with the n-leaf rightist tree. Take the first recursive step in such a recursive construction and replace the whole input stage, as indicated by the reference numeral **7701** in FIG. 77A, with a single $2^n \times 2^n$ node (**7711**) as shown in FIG. 77B. The result is a 2-stage network with a single $2^n \times 2^n$ input node and two $2^{n-1} \times 2^{n-1}$ output nodes. Let the generic step in the recursive construction be changed so that the interconnection is according to this 2-stage network. Unfolding this modified recursive construction, a $2^n \times 2^n$ n-stage unique-routing network results with FIG. 77C depicting the resulting network (**7720**) for n=4.

By filling every node in this network with a 2^k-to-2^{k-1} concentrator with the appropriate k, the resulting n-stage switching network constructs a $2^n \times 2^n$ self-routing nonblocking switch. This n-stage switching network is a hybrid between a routing network and a partial sorting network, to which the self-routing mechanism over the baseline or banyan network, including the priority treatment in Section H, can now be ported over (Note that the guide of the $2^n \times 2^n$ baseline or banyan network is simply the sequence 1, 2, ... , n.) Assuming the presence of the activity bit and the number of priority classes being 2^r, upon entering the switching network, a packet with the binary destination address $d_1 d_2 \ldots d_n$ is preceded by the bit pattern $1 d_1 p_1 \ldots p_r d_2 \ldots d_n$. The in-band

control signal of the packet to a sorting cell is $1d_jp_1...p_r$, which changes only upon the exit of a concentrator. A generic sorting cell inside a concentrator does not alter the in-band control signal carried in the routing tag of the packet. At every output port of the concentrator, a 1×1 element may be installed for the necessary alteration on the routing tag of the packet. When the input traffic has no output contention, such a switching network is a self-routing nonblocking switch. Any packet from any input to any output can be guaranteed.

This construction of a self-routing nonblocking switch can be readily converted into a construction of a self-routing switch that is "nonblocking in the multicast sense", that is, there is no internal blocking whenever there is no output contention among requested connections, point-to-point or multicast. Thus fill every node in the aforementioned n-stage hybrid network with a 2^k-to-2^{k-1} multicast concentrator with the appropriate k. Then, the possible values of an in-band control signal to the concentrator are 0-bound, 1-bound, idle, and bicast. As mentioned in Section H, a multicast concentrator guarantees the maximum possible throughputs at both 0-output group and 1-output group. The multicast concentrator realizes a nonblocking multicast switch when the self-routing multicast mechanism toward an arbitrary set of output addresses as described in the sub-section J1 is ported over. This $2^n \times 2^n$ self-routing switch guarantees the transmission of every packet to all of its destined outputs as long as there is no output contention among the input traffic. In those applications where each packet is destined for a rectangular set of output addresses, the simpler self-routing multicast mechanism as described in Section H can be used instead. This self-routing multicast mechanism for

packets destined for rectangular sets of addresses involves a shorter routing tag and simpler switching control at the network level.

Although the present invention have been shown and described in detail herein, those skilled in the art can readily devise many other varied embodiments that still incorporate these teachings. Thus, the previous description merely illustrates the principles of the invention. It will thus be appreciated that those with ordinary skill in the art will be able to devise various arrangements which, although not explicitly described or shown herein, embody principles of the invention and are included within its spirit and scope. Furthermore, all examples and conditional language recited herein are principally intended expressly to be only for pedagogical purposes to aid the reader in understanding the principles of the invention and the concepts contributed by the inventor to furthering the art, and are to be construed as being without limitation to such specifically recited examples and conditions. Moreover, all statements herein reciting principles, aspects, and embodiments of the invention, as well as specific examples thereof, are intended to encompass both structural and functional equivalents thereof. Additionally, it is intended that such equivalents include both currently known equivalents as well as equivalents developed in the future, that is, any elements developed that perform the function, regardless of structure.

In addition, it will be appreciated by those with ordinary skill in the art that the block diagrams herein represent conceptual views of illustrative circuitry embodying the principles of the invention.--.